

05/19/99



JCS11 U.S. PTO

IN THE U.S. PATENT AND TRADEMARK OFFICE
Express Mail Label No.: EL 260 016 811 US
Attorney Docket No.: T-L Case 4

May 19, 1999

Box Patent Application
Assistant Commissioner for Patents
Washington, D. C. 20231

Sir:

Transmitted herewith for filing is the patent application of:

Inventor(s): Wayne D. GROVER, Demetrios STAMATELAKIS

For : PROTECTION OF ROUTERS IN A
TELECOMMUNICATIONS NETWORK

Enclosed are:

- [X] Specification
[X] Declaration or Oath
[X] 9 Drawing Sheets [X] Formal
[] Informal
[] Preliminary Amendment Cancelling Claims
[] Amendment Before First Office Action
[] Information Disclosure Statement
[] Statement(s) re small entity (37 CFR 1.9 and 1.27)
[X] Assignment, with Recordation Form Cover Sheet
[X] Acknowledgment Postal Card
[] Priority is claimed under 35 USC 119 based on
Application No. , filed .
[] A certified copy of the priority application is enclosed.


The filing fee is calculated as shown below, after entering any
enclosed Amendment which requests entry before such calculation:

For	No. Filed	No. Extra	(X) LG Entity	RATE SM Entity	() Fee
Basic Fee				\$760.00	\$380.00 \$760.00
Total Claims (10 - 20 = 0)			x	\$ 18.00	x \$ 9.00 0.00
Indep. Claims (4 - 3 = 1)			x	\$ 78.00	x \$ 39.00 78.00
[] Multiple Dependent Claim			+	\$260.00	+ \$130.00
* * * TOTAL FILING FEE * * *					\$838.00

[X] A check for \$ 878.00 is enclosed to cover fees.JCS11 U.S. PTO
09/314518
05/19/99

- [] Please charge my Deposit Account No. 06-1382 in the amount of \$ _____. A duplicate of this transmittal is enclosed.
- [X] The Commissioner is hereby authorized to charge payment of the following fees associated with this communication or to credit any overpayment to Deposit Account No. 06-1382. A duplicate copy of this sheet is enclosed.
- [X] Any additional filing fees required under 37 CFR 1.16
- [] Any patent application processing fees under 37 CFR 1.17
- [] Pursuant to 37 CFR 1.52, the enclosed application is in the form of a foreign language text:
- [] An English translation and a statement that the English translation is accurate are enclosed.
- [] Please notify the undersigned of the due date for submitting an English translation and a statement that the English translation is accurate.
- [] An enclosed check includes the \$130.00 fee (37 CFR 1.17k) for processing a foreign language text.

Respectfully submitted,


David S. Goldenberg

FLYNN, THIEL, BOUTELL
& TANIS, P.C.
2026 Rambling Road
Kalamazoo, MI 49008-1699
Phone: (616) 381-1156
Fax : (616) 381-5465

Dale H. Thiel	Reg.No. 24 323
David G. Boutell	Reg.No. 25 072
Ronald J. Tanis	Reg.No. 22 724
Terryence F. Chapman	Reg.No. 32 549
Mark L. Maki	Reg.No. 36 589
David S. Goldenberg	Reg.No. 31 257
Sidney B. Williams, Jr.	Reg.No. 24 949
Timothy B. Clise	Reg.No. 40 957
Liane L. Churney	Reg.No. 40 694
Brian R. Tumm	Reg.No. 36 328

IN DUPLICATE

DSG/T/ad

Encl: Listed Above

100.9810

"Express Mail" Mailing Label No.: EL 260 016 811 US

Date of Mailing: May 19, 1999

Applicant(s) : Wayne D. GROVER, et al.

Title : PROTECTION OF ROUTERS IN A
TELECOMMUNICATIONS NETWORK

Serial No. : Unknown

Filed : Unknown

Atty Docket No.: T-L Case 4

Assistant Commissioner for Patents
Washington, D. C. 20231

EXPRESS MAILING CERTIFICATE

Sir:

I hereby certify that the attached paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Assistant Commissioner for Patents, Washington, D.C. 20231.

FLYNN, THIEL, BOUTELL & TANIS, P.C.

By: Susan J. Proctor

Date: May 19, 1999

Document(s) attached: Patent Application Transmittal dated
May 19, 1999 including enclosures
listed thereon

Telephone: (616) 381-1156
191.9804

TITLE OF THE INVENTION

Protection of Routers in a Telecommunications Network

NAME OF INVENTOR

Wayne D. Grover

Demetrios Stamatelakis

FIELD OF THE INVENTION

This invention relates to methods and node configurations for protecting routers against router failure in a telecommunications network.

BACKGROUND OF THE INVENTION

Internet protocol (IP) networks are a form of packet-based network. Packet are routed in an IP network by the use of routers, which form the nodes of an IP network (node and router can be used interchangeably in this document.) When a router receives an IP packet, it makes the decision of which adjacent router to send the packet to by consulting its local routing tables. The router table entries are indexed by the destination IP address, which is contained in the header of every IP packet. An entry for a destination IP address will generally contain a path cost and a destination address. The path cost is the cost of the previously determined lowest cost path between the router. The destination address is the local router port which will carry the packet to the next router on the path towards the destination address. These routing table entries are determined by the use of a routing protocol (most commonly OSPF in larger IP networks). Thus, an IP packet is generally routed from its source to destination by a series of routers with each router sending the packet closer to its destination by consulting its local routing table. The implication of this is that no high level coordination or setup of the route the packet takes is required; the route the packet takes is determined directly by the sequence of routers it ends up taking.

Traditionally, an IP network's routers restore any failures which take place by execution of the routing protocol to update the router table entries. However, this can take up to a minute with OSPF (and possibly longer with simpler routing protocols) and until the

update takes places the packets which previously would have been routed by a router over/to the failed equipment are simply discarded. This leads to a loss of service quality and availability.

SUMMARY OF THE INVENTION

According to an aspect of the invention, there is provided a method of configuring a node in an IP network by creating a set of router table entries which form a set of virtually preconfigured cyclical routes, or p-cycles, within the IP network. This set of p-cycles are virtual in the sense that the only resources which they normally consume are routing table entries. No transmission capacity is required for them, unless/until they are used to carry packets (traffic.) These p-cycles sit idle and unused until a failure takes place; the routers surrounding the failure then use these p-cycles to route packets, which normally would be lost, around the failure.

According to an aspect of the invention, there is also provided a protected router, comprising a router table, the router table having an entry identifying a cycle of routers encircling an adjacent router to the protected router and a port associated with that entry. The cycle of routers includes all routers logically adjacent to the adjacent router and not the adjacent router.

According to a further aspect of the invention, there is provided a telecommunications network having at least one protected router, and preferably all the routers being protected against failure of adjacent routers.

According to a further aspect of the invention, there is provided a novel data packet comprising an ID field that specifies a p-cycle in which the routers in the p-cycle are all adjacent a router not in the p-cycle, a path cost field and a data field.

These and other aspects of the invention are described in the detailed description of the invention and claimed in the claims that follow.

BRIEF DESCRIPTION OF THE DRAWINGS

There will now be described preferred embodiments of the invention, with reference to the drawings, by way of illustration only and not with the intention of limiting the scope of the invention, in which like numerals denote like elements and in which:

- 5 Figs. 1A-1I show examples of node encircling cycles for respectively nodes 0, 1, 10, 11, 8, 9, 12, 13 and 14 according to the invention;
- Figs. 2A and 2B show re-routing of a packet through a p-cycle;
- Figs. 3A, 3B and 3C show restoration of a node failure using a p-cycle;
- Fig. 4 shows a cycle for a node of degree 2;
- 10 Fig. 5 shows design of a packet according to the invention; and
- Fig. 6 shows a routing table for use in a router according to the invention.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

In this patent document, "comprising" means "including". In addition, a reference to
 15 an element by the indefinite article "a" does not exclude the possibility that more than one of the element is present.

Node Protection of Routers using p-cycles

A p-cycle is found for each potential node failure. Various methods can be used to
 find a p-cycle for each potential node failure including manual and automatic methods. These
 20 p-cycles are to be used to route packets, which normally would be lost, around the failure. This section will focus on the recovery of node failures, but a direct implication of a node failure being directly restorable is that any link (connection between a pair of routers) is also restorable. A p-cycle which can protect the traffic which was routed through a node failure should have two properties: first, it *does not* contain the node which it is protecting, and,
 25 second, it *does* contain all the nodes which are logically adjacent (directly connected) to the protected node. The p-cycle has a hop count lower than any other p-cycle that encircles the node. In general, it may not be possible to always find a simple cycle with these properties (a simple cycle is one which crosses each node, at most once) so cycles which cross a node more than once may also need to be considered.

Examples of simple and non-simple cycles are given in Figs. 1A-1I for various nodes in a 15 node network. Fig. 1A shows an example of a graphically apparent encircling p-cycle for node 0. Fig. 1B shows an example of a graphically non-apparent encircling p-cycle for node 1, in which the p-cycle forms a figure 8. Fig. 1C shows an example of a graphically apparent encircling p-cycle for node 10, which requires inclusion of node 13, not adjacent to node 10, for its construction. Fig. 1D shows an example of a logically encircling p-cycle for node 11, which has a stub, equivalent to a same span figure 8. Fig. 1E shows an example of a graphically non-apparent logically encircling p-cycle for node 8. Fig. 1F shows an example of a graphically non-apparent logically encircling p-cycle for node 9. Fig. 1G shows an example of an encircling p-cycle for node 12. Fig. 1H shows an example of an encircling p-cycle for node 13. Fig. 1I shows an example of an encircling p-cycle for node 14. Thus, it is seen that encircling p-cycles may be visually (graphically) apparent, may require a figure 8, and/or may be non-apparent, that is, not graphically encircling, but still be logically encircling. It is preferred that the set of p-cycles in a network be selected to minimize the number of p-cycles protect all nodes. One way to look at a p-cycle associated with a router is that it should intercept all traffic flows through the router. The node IDs for a p-cycle belonging to a router are stored in all routers in the p-cycle.

A p-cycle, with the properties described in the previous paragraph, can compensate for the loss of the failed node, by offering the lost node's adjacent nodes (which are the only nodes which can route packets to the dead node) an alternate path around the failed node. The cycle covers all the node's adjacent neighbors so a path between them is always possible. So if a node would have routed a packet to a lost node, but detects that the node is now lost, it can divert the packet onto the p-cycle which has been predefined to protect the lost node. The packet then travels the perimeter of the protecting p-cycle, being passed from router to router, until it arrives at a router who "knows" what to do with the packet; at this point the packet is routed normally towards its final destination.

This rerouting of the packet, within the p-cycle, can occur because a router can only route a packet to/from the routers that are immediately adjacent to the router. Therefore, when a router fails, a p-cycle, with the properties described above, can be used to detour packets between pairs of routers that would originally have used the failed router. The p-

cycle is guaranteed to provide this detour because it is designed to cover all the failed router's adjacent neighbors and to not contain the failure node. An example of this packet rerouting is given in Figs. 2A and 2B. In Fig. 2A, if the router 10 fails, a packet being sent from source node 12 to destination node 14 is routed as shown in Fig. 2B around nodes 16, 18, 20 and 22 which form part of a p-cycle surrounding node 10. The complete cycle also includes nodes 12, 14 and 24.

Because a p-cycle can only be formed as a virtual circuit within a packet switched IP network, a p-cycle does not use any transmission capacity unless it is actually used. For example, in node protection it is conceivable that a separate p-cycle will be formed for each network router. However, in normal operation these cycles will not use any transmission capacity. Only when a node in the network fails would a p-cycle utilize capacity to detour packets around the failure. This is in contrast to a circuit switched network where protection capacity used to form a p-cycle would use up the capacity regardless of if it is used to restore a failure.

Use of P-cycles to Restore Failures

As previously discussed a p-cycle restores a failure by providing a detour for a packet around a node failure. The packet is rerouted through the p-cycle until it reaches a router that is able to remove the packet from the p-cycle and route it normally. Some simple rules are required, to decide if a router may safely remove a packet from a p-cycle, mainly to prevent the formation of packet loops. A packet travels in a loop if somehow it repeatedly enters and exits a p-cycle (a packet should enter and exit a p-cycle only once.)

The rule used to prevent these loops is based on path cost. Each routing table entry for a destination IP address has a field containing the path cost from the local router to the destination. When being normally routed, the path cost of a packet's routing table entry will go down as the packet gets closer to its destination because, by design, a packet is routed towards its destination along lowest cost route. This would not be true if the cost were to increase as a packet progressed to its destination.

To prevent looping within the p-cycle care must be taken that a packet exits the p-cycle at a point that is "downhill" from the point where it entered ("downhill" in the sense that the exit point is closer to the packet's destination than the entry point.) This prevents

loops because for a packet to return to the p-cycle it would effectively have to move further away from its destination; this is not possible if the packet is routed normally using lowest cost paths.

P-cycle processing rules at the insertion router:

5 If a router attempts to send a packet through a port but cannot because the router on the other side has been determined to be dead, the router encapsulates the original IP packet within a p-cycle packet and sets the path cost field within the p-cycle packet to equal the cost of the path the packet would have taken had the router failure not occurred. The structure of the packet is shown in Fig. 5, with p-cycle ID field 50 and path cost field 52. The p-cycle ID
10 field 50 contains the ID of the p-cycle on which the packet belongs. The path cost field 52 contains the path cost of the original pre-failure path for the IP packet. The router then accesses the label of the p-cycle that was dedicated to protecting the failed router and inserts the p-cycle packet into the p-cycle.

P-cycle processing rules at subsequent router on the p-cycle:

15 Subsequent routers, upon reception of a p-cycle packet, determine the destination address of the encapsulated IP packet and use this to access the local routing table at the subsequent router. If the entry in the local routing table for the destination node points to a dead port (presumably, due to the router failure), or there is no entry for the destination node, the router continues relaying the p-cycle packet along the p-cycle. If the port is not
20 dead, the router compares the local table entry's path cost to the p-cycle packet's cost field. If the local cost is greater than or equal to packet's cost the packet is allowed to continue along the p-cycle. However, if the local cost is less than the packet's cost, the original IP packet is retrieved from the p-cycle packet and routed normally (the router's location is suitably "downhill" from the p-cycle packet's entry point, and it is safe to route the IP packet
25 normally.)

As previously, mentioned it is guaranteed that eventually the packet will arrive at a router which can remove the packet from the p-cycle. However, to safeguard against the situation where no route off the p-cycle exists (perhaps another failure occurred, eliminating any possible route), an additional check can be introduced where each router checks if it had
30 originally introduced the packet to the p-cycle so it can dispose of the packet if it was. This

check prevents a packet from continually looping around the p-cycle. Note that this check would imply the presence of a field in a p-cycle packet which would contain the ID of the router which introduced the packet into the p-cycle.

Fig. 3 gives an example of these p-cycle routing rules. In the example, a packet would originally have been traveled over the route from node 12 to node 14 given in Figure 3a. However, a node failure at node 10 disrupts this route and forces the source router 12 to detour the packet onto a p-cycle, after encapsulating it in a p-cycle packet (Figure 3b.) The first router 16 along the p-cycle continues to pass the packet along as the path cost it has for the packet destination is infinite (its route for the packet has also been disrupted by the failure.) The second router 18 does have a valid route for the packet but it also passes the packet along the p-cycle. The local path cost of router 18 is 3 which is equal to that of the original route and so the packet can not be safely handled using conventional routing. At the third router 20, however, the local path cost is 2 (less than the original cost of 3) and, so, the packet is routed normally to its destination node 14 (Fig. 3c).

An exemplary routing table is shown in Fig. 6. The routing table has a set of addresses, including destination addresses, and also includes a p-cycle address. For each address, there is a port ID entry, and a cost for using that port to get to the destination. The p-cycle address specifies a port leading to the next router in the p-cycle. Fig. 6 also shows schematically that if an IP packet arrives with a destination for which the corresponding port is dead, the data packet is encapsulated as a p-cycle packet and routed according to the p-cycle address entry and its corresponding port.

Considerations in Node restorable p-cycle design

Two properties are required for a node to be protected by a p-cycle. The p-cycle must cover all adjacent nodes (to ensure routes between the protected node and adjacent neighbors). The p-cycle must not contain the node to be protected (so that the p-cycle is not disrupted when the node fails.) This can, however, raise issues when dealing with adjacent nodes that are of degree 2 (only 2 links terminate on the node.) The complication that arises is that the failure of a node connected to a degree 2 node will partially disconnect the node resulting in only a single access point. This problem with using p-cycles with such a site is p-

cycles can only be formed with nodes that are of at least degree 2 *after* the node failure occurs.

This can be handled by building a p-cycle for a node failure by considering only the adjacent nodes which are of degree 3 or higher. The nodes which are of degree 2 can then be
5 connected into the p-cycle using a point to point segment. This is a special case of a figure 8 loop in which the figure 8 loop is logically required to go down and back the same span, to include one or more degree two sites. Fig. 4 gives an example of such an arrangement. Nodes 30, 32, 34, 36 and 38 are of at least degree 3. Nodes 40 and 42 are of degree 2. Faced
10 with a failure a degree 2 node such as node 40 encapsulates an IP packet as usual and introduces it into its part of the preconfigured segment and forwarded to node 38. From there it would be passed to the main part of the p-cycle where it would be passed along the perimeter using the normal p-cycle rules. The segment connecting nodes 38 and 40 acts as a one way entrance into the p-cycle. Once within the p-cycle a packet will only rotate around the cycle; it will not exit using the segment. Packets on the p-cycle which need to reach a
15 degree 2 node can do this using the previously described mechanism of removal from the p-cycle from which point they will be normally routed towards the node.

Immaterial modifications may be made to the invention described here without departing from the essence of the invention.

THE EMBODIMENTS OF THE INVENTION IN WHICH AN EXCLUSIVE
PROPERTY OR PRIVILEGE IS CLAIMED ARE DEFINED AS FOLLOWS:

- 5 1. A telecommunications network, comprising:
plural routers; and
at least one protected router comprising a router table, the router table having an
entry identifying an alternative route around an adjacent router to the protected router in case
of failure of the adjacent router and a port associated with that entry.
- 10 2. The telecommunications network of claim 1 in which there is associated with each
router in the alternative route has a routing table with an entry identifying the alternative
route and a port associated with that entry.
- 15 3. A protected router, comprising a router table, the router table having an entry
identifying a cycle of routers encircling an adjacent router to the protected router and a port
associated with that entry, the cycle of routers including all routers logically adjacent to the
adjacent router and not the adjacent router.
- 20 4. The protected router of claim 3 in which the cycle of routers includes only logically
adjacent routers to the protected router.
- 25 5. The protected router of claim 3 in which the protected router has a router table in
which is stored, for each adjacent router to the protected router, an entry identifying a cycle
of routers encircling an adjacent router to the protected router and a port associated with that
entry, each cycle of routers including all routers logically adjacent to the respective adjacent
router and not the respective adjacent router.
6. A data packet comprising:

an ID field that specifies a p-cycle of routers in which the routers in the p-cycle are all adjacent a router not in the p-cycle, a path cost field and a data field.

7. A method of protecting against router failure in a network, in which the network includes plural interconnected routers, the method comprising the step of:

storing at a protected router an entry identifying a cycle of routers that form at least one alternative route around an adjacent router to the protected router, in which the cycle of routers includes all routers logically adjacent to the adjacent router and not the adjacent router.

8. The method of claim 7 further comprising the step of:

upon failure of the adjacent router, routing all data packets whose least cost path includes the adjacent router, around the alternative route beginning at the protected router.

9. The method of claim 8 in which each data packet routed around the alternative route contain an ID field that identifies the cycle of routers, a path cost field containing the cost of the least cost path and a data field.

10. The method of claim 9 in which each router in the alternative route has a router table having an entry that identifies the cycle of routers and continues to route the data packet around the alternative route until the path cost from a router in the alternative route to the destination of the data packet is less than the cost of the least cost path.

ABSTRACT OF THE DISCLOSURE

A method of configuring a node in an IP network by creating a set of router table entries which form a set of virtually preconfigured cyclical routes, or p-cycles, within the IP network. This set of p-cycles is virtual in the sense that the only resources which they normally consume are routing table entries. No transmission capacity is required for them, unless/until they are used to carry packets (traffic.) These p-cycles sit idle and unused until a failure takes place. The routers surrounding the failure then use these p-cycles to route packets, which normally would be lost, around the failure. A router having an entry in its router table identifying the p-cycle, together with an associated port, is also disclosed, and also a data packet that may use a p-cycle to get to its destination.




FIGURE 1B




FIGURE 1A




FIGURE 1E




FIGURE 1F




FIGURE 1H




FIGURE 1G




FIGURE 11




FIGURE 2A




FIGURE 2B




FIGURE 3A




FIGURE 3B




FIGURE 3C




FIGURE 4




FIGURE 5

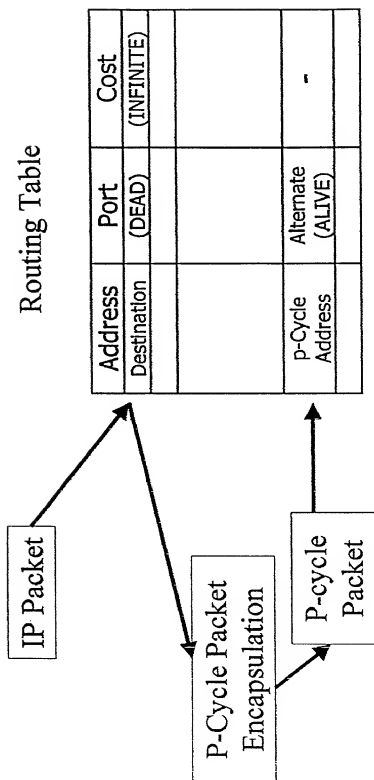


FIGURE 6

DECLARATION FOR PATENT APPLICATION

Page 1 of 2

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name.

I believe I am the original, first and sole inventor (if one name is listed below) or an original, first and joint inventor (if more than one name is listed below) of the subject matter which is disclosed and/or claimed and for which a patent is sought on the invention entitled: **Protection of Routers in a Telecommunications Network** the specification of which

xxx is attached hereto (as amended on _____)
_____ WAS FILED _____ as Application Serial No. _____
and was amended on _____.

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose to the Patent Office all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, §1.56(a).

I hereby appoint the following attorneys and/or agents to prosecute this application and to transact all business in the United States Patent and Trademark Office connected therewith: Dale H. Thiel, Reg. No. 24,323; David G. Boutell, Reg. No. 25,072; Ronald J. Tanis, Reg. No. 22,724; Terryence F. Chapman, Reg. No. 32,549; Mark L. Maki, Reg. No. 36,589; David S. Goldenberg, Reg. No. 31,257; Sidney B. Williams, Jr. Reg. No. 24,949; Timothy B. Clise, Reg. No. 40,957; Liane L. Churney, Reg. No. 40,694; Brian R. Tumm, Reg. No. 36,328; and the firm of Flynn, Thiel, Boutell & Tanis, P.C. and Anthony R. Lambert, Reg. No. 32,813.

Address all telephone calls to David S. Goldenberg at telephone No. (616)381-1156.

Address all correspondence to:

FLYNN, THIEL, BOUTELL & TANIS, P.C.
2026 Rambling Road
Kalamazoo, MI 49008-1699

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true;

Declaration for application

Page 2 of 2

and further that these statements were made with the knowledge that wilful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such wilful false statements may jeopardize the validity of the application or any patent issued thereon.

Wayne D. Grover

FULL NAME OF SOLE/FIRST INVENTOR

W.D. Grover

INVENTOR'S SIGNATURE

May 17/99

DATE

Edmonton, Alberta, Canada

RESIDENCE

Canadian

CITIZENSHIP

8944 - 116 St., Edmonton, Alberta, Canada T6G 1P8

POST OFFICE ADDRESS

Demetrios Stamatelakis

FULL NAME OF JOINT/SECOND INVENTOR

Demet St

INVENTOR'S SIGNATURE

May 14, 1999

DATE

Edmonton, Alberta, Canada

RESIDENCE

Canadian

CITIZENSHIP

15131 - 43 Avenue, Edmonton, Alberta, Canada, T6H 5L7

POST OFFICE ADDRESS